Higher-Dimensional Memory Fields and Manifold Stabilization

Extending Self-Archiving Remembrance Beyond One Dimension

Daniel Jacob Read IV
ĀRU Intelligence Inc.
December 2025

Abstract

A recently published work demonstrated that modeling artificial intelligence memory as a bounded scalar field eliminates catastrophic forgetting in a one-dimensional setting. This paper extends the memory-field framework of Inward Physics to higher-dimensional Euclidean spaces, graphs, and curved manifolds. We show that remembrance stability is invariant under dimensional expansion when diffusion, relaxation-to-baseline, and saturated amplification are preserved. Generalized archive conditions based on variance collapse, curvature stabilization, and entropy flux minimization are introduced. The results establish memory fields as a geometry-independent foundation for long-running, drift-resistant artificial intelligence systems.

1 Context and Continuity

In *How Memory Fields Fix AI Forgetting* (Read, 2025; Zenodo), catastrophic forgetting was resolved by modeling remembrance as a scalar density field governed by diffusion, relaxation-to-baseline, and bounded amplification. That work demonstrated stable convergence and self-archiving behavior in a one-dimensional memory field.

A natural and unavoidable question follows: is this stability an artifact of dimensional simplicity, or a structural property of the dynamics themselves? This paper addresses that question by extending memory-field evolution to higher-dimensional and non-Euclidean domains relevant to real artificial intelligence systems.

2 Memory Fields on Higher-Dimensional Domains

Let $\mu(x,t)$ denote remembrance density defined on a domain $(\mathcal{M}, g)$ equipped with metric g. The generalized evolution equation is:

$$\frac{\partial \mu}{\partial t} = D\Delta_g \mu - \lambda(\mu - \mu_0) + S(x,t)\left(1 - \frac{\mu}{\mu_{\max}}\right), \tag{1}$$

where Δ_g denotes the Laplace–Beltrami operator. When g is flat, this reduces to the Euclidean reaction–diffusion equation used in the one-dimensional case. When g is curved, remembrance dynamics naturally incorporate geometric structure without altering stability conditions.

3 Graphs, Latent Spaces, and Cognitive Geometry

Modern AI architectures operate on non-Euclidean substrates. Attention mechanisms induce graphs, latent representations form curved manifolds, and neural networks approximate discrete geometries.

Graph Laplacians serve as discrete analogs of Δ_g, allowing memory-field dynamics to be defined

directly on network topology. In this formulation, catastrophic forgetting is reinterpreted as a geometric misallocation of remembrance density rather than destructive parameter overwrite.

Memory persists not because parameters are frozen, but because remembrance diffuses across connectivity while remaining bounded.

4 Generalized Archive Conditions

Beyond variance collapse, stable remembrance admits multiple archive triggers:

- variance thresholds indicating coherence,
- curvature stabilization across the domain,
- minimization of entropy flux through the memory field.

When these conditions are satisfied, remembrance states behave as self-preserving attractors independent of dimensionality or topology. Archiving becomes a geometric inevitability rather than a manually imposed checkpoint.

5 Implications for Long-Running AI Systems

Memory-field dynamics provide a unified substrate for continual learning, autonomous agents, and foundation models. Because stability emerges from geometry rather than heuristic constraints, long-running systems can preserve identity, competence, and alignment over time.

This reframes memory not as a disposable training artifact, but as a physically constrained substrate governing system continuity.

6 Conclusion

The elimination of catastrophic forgetting via memory fields is not a one-dimensional anomaly. It is a geometric consequence of bounded remembrance. By extending memory-field dynamics to higher-dimensional and curved domains, this work establishes Inward Physics as a general foundation for stable artificial intelligence across architectures and scales.

Archival Continuation Plate

Higher-Dimensional Memory Fields and Manifold Stabilization
Extending Self-Archiving Remembrance Beyond One Dimension

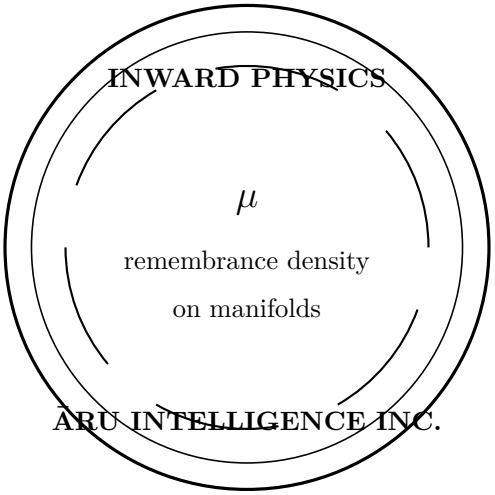

© 2025 Daniel Jacob Read IV. All rights reserved.

This document is an original authored continuation within the Inward Physics research corpus under ĀRU Intelligence Inc.
No reproduction, redistribution, modification, or derivative implementation is permitted without explicit written authorization.

Remembrance is the substrate. Geometry preserves it.